arXiv:1606.02386v1 [stat.ME] 8 Jun 2016

A Generalized Benjamini-Hochberg Procedure for Multivariate Hypothesis Testing

Kasra Alishahi*, Ahmad Reza Ehyaei*, and Ali Shojaie**

*Department of Mathematical Sciences, Sharif University of Technology, Tehran, Iran
**Department of Biostatistics, University of Washington, Seattle, WA

June 9, 2016

Abstract

The introduction of the false discovery rate (FDR) by Benjamini and Hochberg has spurred a great interest in developing methodologies to control the FDR in various settings. The majority of existing approaches, however, address the FDR control for the case where an appropriate univariate test statistic is available. Modern hypothesis testing and data integration applications, on the other hand, routinely involve multivariate test statistics. The goal, in such settings, is to combine the evidence for each hypothesis and achieve greater power, while controlling the number of false discoveries. This paper considers data-adaptive methods for constructing nested rejection regions based on multivariate test statistics (z-values). It is proved that the FDR can be controlled for appropriately constructed rejection regions, even when the regions depend on data and are hence random. This flexibility is then exploited to develop optimal multiple comparison procedures in higher dimensions, where the distribution of non-null z-values is unknown. Results are illustrated using simulated and real data.

Keywords: False discovery rate, Multivariate z-value, Random nested rejection regions, High-dimensional statistics

1 Introduction

Multiple hypothesis testing is a fundamental problem in many new scientific applications involving Big Data, from genomics and neuroscience to astronomy and finance. The false discovery rate (FDR) control, introduced in

Benjamini and Hochberg's seminal paper [5], is one of the most important methodological developments in multiple hypothesis testing. To control the FDR at a predetermined level α, Benjamini and Hochberg (BH) proposed a step-down procedure based on ranked p-values. The procedure was initially developed under the assumption that the n_0 null hypotheses are independent of each other and are also independent of the n_1 non-null hypotheses; nonetheless, it continues to control the FDR if the $n = n_0 + n_1$ hypotheses are positively dependent; see, e.g., Benjamini and Yekutieli [7] for additional details.

A potential drawback of the BH procedure is that it controls the FDR at the level $(n_0/n)\alpha$ and is hence conservative. A number of authors, including, Benjamini and Hochberg [6], Storey et al. [20], Benjamini et al. [8] and Gavrilov et al. [13], have thus proposed modifications of the BH procedure for more efficient multiple testing by estimating n_0. These 'BH-type' procedures primarily focus on better control of the FDR, but their efficiency has not been formally investigated. In particular, while they guarantee the control of FDR $= \mathbb{E}(V/R)$, for V and R defined in Table 1, they do not provide any guarantees on the false negative rate FNR $= \mathbb{E}\left(T/(n-R)\right)$.

As an alternative to controlling the FDR based on p-values, a number of authors have advocated the use of test statistics, or z-values. Efron [12] introduced the control of *local* false discovery rate (Lfdr), which facilitates the calculation of size and power in large-scale testing problems. Consider a two component mixture model $f(z) = \pi_0 f_0(z) + \pi_1 f_1(z)$, in which f_0 and f_1 are densities of z-values under the null and non-null hypotheses; for each $1 \leq i \leq n$, $H_{0,i}$ is then true with probability π_0 and false with probability $\pi_1 = 1 - \pi_0$. For z-values $Z_1, Z_2, \ldots, Z_n$ corresponding to the $H_{0,i}$, Efron defined the Lfdr as

$$\text{Lfdr}(z) = \mathbb{P}\left(H_{0,i} \text{ is true } \mid Z_i = z\right) = \frac{\pi_0 f_0(z)}{f(z)},$$

and proposed to reject $H_{i,0}$ whenever Lfdr(Z_i) does not exceed a threshold $\lambda > 0$.

Table 1: Possible outcomes in multiple hypothesis testing problems.

	Decisions		
Truth	Not rejected	Rejected	Total
Null	U	V	n_0
Non-Null	T	S	n_1
Total	$n - R$	R	n

As yet another alternative to the FDR, Sun and Cai [21] and Xie et al. [23] developed adaptive multiple testing procedures based on z-values and showed that their procedure is *optimal* in the sense that it minimizes the empirical false negative rate, $\text{mFNR} = \mathbb{E}(T)/\mathbb{E}(n-R)$, while controlling the empirical false discovery rate, $\text{mFDR} = \mathbb{E}(V)/\mathbb{E}(R)$; see Table 1. Let $F_0(z)$, $F_1(z)$ and $F(z)$ be cumulative density functions (cdf's) corresponding to $f_0(z)$, $f_1(z)$ and $f(z)$, respectively. Alternatively, mFDR can then be defined as the posterior probability of a case being null given that its z-value Z_i is less than some cutoff z,

$$\text{mFDR}(z) = \mathbb{P}(H_{0,i} \text{ is true } \mid Z_i \leq z) = \frac{\pi_0 F_0(z)}{F(z)}.$$

It can thus be seen that mFDR and Lfdr are analytically related:

$$\text{mFDR}(z) = \int_{-\infty}^{z} \text{Lfdr}(u) f(u) du \Big/ \int_{-\infty}^{z} f(u) du = \mathbb{E}\left[\text{Lfdr}(Z) \mid Z \leq z\right].$$

Genovese and Wasserman [14] showed that for independent hypotheses, mFDR (mFNR) and FDR (FNR) are asymptotically equivalent, in the sense that $\text{mFDR} = \text{FDR} + \mathcal{O}(n^{-1/2})$. Thus, the procedure of Sun and Cai asymptotically controls the FDR. However, procedures based on z-values may not provide exact FDR control at a given level α.

Despite significant progress in multiple hypothesis testing, existing approaches are mainly suitable for univariate hypotheses, i.e., for testing multiple univariate p- or z-values. Increasingly, however, multivariate, and potentially high-dimensional, evidence is available for each hypothesis. A prime example of such applications arises in *data integration* for high-throughput biology: As new high-throughput technologies emerge, biomedical scientists now collect multiple types of data on various aspects of cellular function, including DNA variants, copy number variation, DNA methylation, mRNA expression and abundances of proteins and metabolites. The Cancer Genome Atlas (TCGA) [3], the ENCODE project [1] and the Genotype-Tissue Expression (GTEx) project [2] are just a few examples of massive efforts to collect diverse high-throughput data in order to accelerate scientific discoveries. The hope, in these and other projects, is to delineate cellular functions and mechanisms of disease initiation and progression by integrating the evidence from diverse high-throughput data.

A key step in ensuring the reproducibility of findings from testing many multivariate hypotheses is controlling the number of false discoveries. However, when $d > 1$, there is no unique way to rank multivariate hypotheses

Table 2: Estimated FDR and FNR for method of Sun and Cai [21] (SC) and the oracle method in high dimensions. As the dimension d increases, SC fails to control the FDR at the pre-specified level of 10%.

d	SC FDR	Oracle FDR	SC FNR	Oracle FNR
2	0.11	0.10	0.13	0.13
3	0.14	0.10	0.12	0.12
4	0.20	0.10	0.09	0.06
5	0.40	0.10	0.09	0.03
6	0.52	0.10	0.07	0.02
7	0.66	0.10	0.08	0.02
8	0.73	0.10	0.08	0.07
9	0.76	0.10	0.04	0.00
10	0.79	0.10	0.04	0.00
11	0.80	0.10	0.01	0.01
12	0.80	0.10	0.01	0.00
13	0.80	0.10	0.00	0.00
14	0.80	0.10	0.01	0.00
15	0.80	0.10	0.05	0.00

tests based on p-values, which is necessary for applying BH-type procedures. It is therefore not clear how to use BH-type procedures in order to control the FDR in multivariate settings. On the other hand, methods based on mixture model and density estimation can, in principal, be applied when $d > 1$. An early attempt to control the FDR in the multivariate setting was the proposal of Ploner et al. [18], who tried to generalize Efron's Lfdr as a function of multivariate z-values. They proposed to estimate f_0, f_1 and π_0 and then construct a local FDR function for rejecting $d = 2$ dimensional hypotheses. However, Efron's method assumes that Lfdr(z) decreases as $|z|$ becomes larger. Under this assumption, FDR(z) is smaller than Lfdr(z), so controlling the Lfdr also controls the FDR. For instance, if $f_{(i)}$ is the standard normal density, then the family of densities $\{f_{(i)}(z-\mu) : \mu \in \mathbb{R}\}$ has the monotone decreasing local FDR property (see, e.g., [21] for more details). Therefore, in this case, controlling the Lfdr(z) at any level guarantees that the FDR does not exceed that threshold. However, in many cases, for instance when $f_{(i)}$ have heavy tails or multiple modes, Lfdr(z) is no longer monotone decreasing. In such settings, the methods based on Lfdr, e.g. [12] and [18], are not guaranteed to control the FDR.

Another option for multivariate FDR control is the proposal of Sun and Cai

(SC), which requires estimation of the multivariate densities, f_0 and f_1, as well as the proportion of null hypotheses π_0. Unfortunately, density estimation becomes increasingly difficult as the dimension d increases. For instance, multivariate kernel density estimation is challenging when $d > 5$, the best possible (minimax) rate of mean-squared error of kernel density estimation is $\mathcal{O}(n^{-4/(4+d)})$ [19]. This bound underscores the "curse of dimensionality" when d is large. As a result, the asymptotic FDR control of Sun and Cai's method becomes invalid when d is large. Table 2 illustrates the inability of the SC method to control the FDR in multivariate settings in comparison to an oracle procedure, which assumes that f_1 and π_0 are known. Here, $n = 10\,000$ z-values are generated from a mixture of Gaussians with $\pi_0 = 0.8$; the null hypotheses are independent standard normals, and the alternative hypotheses have a Gaussian density with a random covariance matrix and mean vector of length one. Since estimating π_0 in high dimensions is not straightforward (for instance the method of Jin and Cai [15] is not applicable), in this simulation, we have used the true value of $\pi_0 = 0.8$. The results clearly show that even for moderate dimensions, e.g. when $d = 3$ or 4, SC fails to control the FDR at the nominal level of $\alpha = 0.10$.

An alternative to FDR control for multivariate hypotheses is to combine the evidence from the d-variate hypotheses into a single summary measure. For instance, one can apply the Fisher's transformation for combining the d p-values corresponding to each test of hypothesis to obtain a single p-value. BH-type methods can then be used to control the FDR using the resulting univariate p-value. However, the validity of such summaries often relies on strong assumptions. Moreover, as we will show in Section 5.1 such approaches can also be inefficient.

In this paper, we propose a new procedure that overcomes the shortcomings of existing approaches for controlling the FDR in multivariate settings. The proposed procedure utilizes a stepwise rejection strategy to achieve exact FDR control, and is thus a BH-type procedure. However, our approach is based on test statistics, or z-values. To achieve FDR control for arbitrary d, at each step of the procedure, we use not the entire data, but the part corresponding to previously rejected hypotheses. In other words, once a hypothesis is rejected, we use its z-value to improve our estimate of the non-null distribution, and consequently the likelihood ratio statistics for other hypotheses. To show that the proposed procedure achieves exact control of the FDR, we first present a new proof for the BH step-down procedure. The new proof uses techniques from stochastic calculus and is amenable to higher dimensions. It also allows us to use random data-driven rejection regions. The key to controlling the FDR using this new procedure is that,

in each step, we are not allowed to change our decision about previously rejected hypotheses, nor to use the z-values corresponding to hypotheses that are not yet rejected. Fortunately, since hypotheses are rejected according to their (estimated) likelihood ratio statistics, the previously rejected z-values contain the most information about the non-null distribution. We are thus able to achieve asymptotically optimal multiple hypothesis testing in high dimensions. Throughout the paper, we formulate our procedure and results in terms of more geometric notions of rejection regions rather than the ultimate univariate test statistic. However, these formulations become equivalent when the regions are level sets of d-dimensional test statistics.

The rest of the paper is organized as follows. In Section 2, we present our new proof of the BH step-down procedure. Using this result, in Section 3 we establish exact FDR control in higher dimensions. In Section 4, we first present an oracle decision rule for multivariate hypothesis testing assuming that the non-null distribution f_1 is known. The generalization of this approach to the settings where f_1 is estimated is presented Section 4.1. The results of applying the proposed procedure to simulated and real data examples, as well as comparisons to existing approaches are presented in Sections 5.1 and 5.2, respectively. We conclude the paper with a discussion in Section 6.

2 A New Proof for Step-Down BH Procedure

Let $H_1, \ldots, H_n$ be n null hypotheses and $p_1, \ldots, p_n$ be their corresponding p-values with $p_i \sim \text{Unif}[0,1]$ whenever the ith null hypothesis is true. We assume that p_i's corresponding to null hypotheses are independent of each other and also independent of the p-values of non-null hypotheses. Let $p_{(1)} \leq \ldots \leq p_{(n)}$ be the ordered p-values, and $H_{(1)}, \ldots, H_{(n)}$ be their corresponding hypotheses.

Consider the following empirical processes

$$\begin{aligned} a_t &= \#\{\text{null } p_i : p_i \leq t\}, \\ b_t &= \#\{\text{non-null } p_i : p_i \leq t\}, \\ r_t &= a_t + b_t. \end{aligned}$$

The false discovery proportion (FDP) and false discovery rate (FDR) can

then be defined as

$$\mathrm{FDP}(t) = \frac{a_t}{r_t \vee 1},$$
$$\mathrm{FDR}(t) = \mathbb{E}\left[\frac{a_t}{r_t \vee 1}\right].$$

The step-down BH procedure rejects the hypotheses $H_{(1)}, \ldots, H_{(i_{sd})}$, where

$$i_{sd} = \max\{i : p_{(i)} \leq \tau_{sd}\}, \text{ with } \tau_{sd} = \inf\left\{t : \ r_t \leq \frac{nt}{q}\right\}, \tag{1}$$

and accepts the rest. Here q is a user-specified "error" rate, or more precisely, the desired level of FDR control. Benjamini and Hochberg [5] showed that the step-down BH procedure controls the FDR at the level q.

We next present an alternative proof for the step-down BH procedure, using tools from stochastic calculus. Using the new proof, in Section 3 we generalize the BH procedure to multivariate hypotheses. An important feature of the new proof is that it allows the rejection region to be random and to vary depending on the previously rejected hypotheses. This property will prove particularly useful in the development of FDR controlling procedures for multivariate hypothesis testing in Sections 3 and the proposed optimal procedure in Sections 4.

Theorem 1. *If the p-values corresponding to null hypotheses are independent of other p-values, then the step-down BH algorithm* (1)*, denoted* $BH_{sd}(qr_{\cdot})$*, controls the FDR at the level* q*, i.e.,*

$$\mathbb{E}\left[\frac{a_{\tau_{sd}}}{r_{\tau_{sd}}}\right] \leq \pi_0 q \leq q. \tag{2}$$

Proof. By the definition of τ_{sd} in (1), the statement (2) is equivalent to

$$\mathbb{E}\left[\frac{a_{\tau_{sd}}}{r_{\tau_{sd}}}\right] = \mathbb{E}\left[\frac{q}{n}\frac{a_{\tau_{sd}}}{\tau_{sd}}\right] \leq q.$$

However, $r_t = a_t + b_t \leq \dfrac{nt}{q}$, and thus, $a_t \leq \dfrac{nt}{q} - b_t$. Given that b_t is independent of a_t, it suffices to show that $\mathbb{E}\left[\dfrac{a_{\tau_{sd}}}{r_{\tau_{sd}}} \mid b_t\right] \leq q$. We can thus assume, without loss of generality, that b_t is an arbitrary deterministic process.

Now, let $\hat{a}_t = a_t - n_0 t$, where n_0 is the number of null hypotheses. Then, τ_{sd} can be redefined as

$$\tau_{sd} = \inf\left\{t : \ \hat{a}_t \leq \frac{nt}{q} - b_t - n_0 t\right\}.$$

We can then write

$$\mathbb{E}\left[\frac{a_{\tau_{sd}}}{r_{\tau_{sd}}}\right] = \mathbb{E}\left[\frac{q}{n}\frac{a_{\tau_{sd}}}{\tau_{sd}}\right] \leq \mathbb{E}\left[\frac{q}{n_0}\frac{a_{\tau_{sd}}}{\tau_{sd}}\right] \leq q.$$

Equivalently,

$$\mathbb{E}\left[\frac{1}{n_0}\frac{a_{\tau_{sd}}}{\tau_{sd}}\right] \leq 1, \text{ or, } \mathbb{E}\left[\frac{\hat{a}_{\tau_{sd}}}{\tau_{sd}}\right] \leq 0.$$

The statement of the theorem now follows by the following property of the process $\hat{a}_t$, established in Lemma 2. □

Lemma 2. *Let $\hat{a}_t$ be a process as defined above and g be an arbitrary deterministic function. For a stopping time τ defined as*

$$\tau = \inf\{t : \hat{a}_t \leq g(t)\},$$

we have

$$\mathbb{E}\left[\frac{\hat{a}_\tau}{\tau}\right] \leq 0. \tag{3}$$

Proof of Lemma 2. We first prove that $\mathbb{E}\left[\frac{\hat{a}_\tau}{\tau+\epsilon}\right] \leq 0$. Note that

$$\frac{\hat{a}_\tau}{\tau+\epsilon} = \int_0^\tau \frac{1}{s+\epsilon}\, d\hat{a}_s - \int_0^\tau \frac{\hat{a}_s}{(s+\epsilon)^2}\, ds. \tag{4}$$

Now, we claim that the process η_t defined as

$$d\eta_s = d\hat{a}_s + \frac{\hat{a}_s}{1-s}\, ds, \tag{5}$$

is a martingale with respect to the filtration $\mathcal{F}_s = \sigma(a_u, u \leq s)$. To prove this claim, it suffices to show that $\mathbb{E}[\eta_t - \eta_s \mid \mathcal{F}_s] = 0$. But,

$$\begin{aligned}\mathbb{E}[\eta_t - \eta_s \mid \mathcal{F}_s] &= \mathbb{E}\left[\hat{a}_t - \hat{a}_s + \int_s^t \frac{\hat{a}_u}{1-u}\, du \mid \mathcal{F}_s\right] \\ &= \mathbb{E}[\hat{a}_t - \hat{a}_s \mid \mathcal{F}_s] + \int_s^t \frac{\mathbb{E}[\hat{a}_u \mid \mathcal{F}_s]}{1-u}\, du.\end{aligned} \tag{6}$$

It is then easy to see that the process $\hat{a}_t$ satisfies

$$\mathbb{E}[\hat{a}_t \mid \mathcal{F}_s] = \frac{1-t}{1-s}\hat{a}_s,$$

which is true because $a_t - a_s$ is independent of $\mathcal{F}_s$ and has binomial distribution $B\left(n_0 - a_s, \frac{t-s}{1-s}\right)$. The right-hand side of Equation (6) can thus be written as

$$\mathbb{E}[\hat{a}_t|\mathcal{F}_s] = \frac{s-t}{1-s}\hat{a}_s + \int_s^t \frac{\frac{1-u}{1-s}\hat{a}_s}{1-u} du = 0.$$

Now, substituting Equation (5) into Equation (4), we get

$$\begin{aligned}\frac{\hat{a}_\tau}{\tau+\epsilon} &= \int_0^\tau \frac{1}{s+\epsilon}\, d\eta_s - \int_0^\tau \frac{\hat{a}_s}{(s+\epsilon)(1-s)}\, ds - \int_0^\tau \frac{\hat{a}_s}{(s+\epsilon)^2}\, ds \\ &= \int_0^\tau \frac{1}{s+\epsilon}\, d\eta_s - \int_0^\tau \frac{1+\epsilon}{(s+\epsilon)^2(1-s)}\hat{a}_s\, ds.\end{aligned}$$

Taking expectation, we obtain:

$$\mathbb{E}\left[\frac{\hat{a}_\tau}{\tau+\epsilon}\right] = \mathbb{E}\left[\int_0^\tau \frac{1}{s+\epsilon}\, d\eta_s\right] - \mathbb{E}\left[\int_0^\tau \frac{1+\epsilon}{(s+\epsilon)^2(1-s)}\hat{a}_s\, ds\right]. \qquad (7)$$

The first term in Equation (7) is an integral with respect to a martingale and hence a martingale. Also, τ is a stopping time. So, by the optional sampling theorem [17], the expectation vanishes. Thus,

$$E\left[\frac{\hat{a}_\tau}{\tau+\epsilon}\right] = -E\left[\int_0^\tau \frac{1+\epsilon}{(s+\epsilon)^2(1-s)}\hat{a}_s\, ds\right].$$

From this we have

$$\begin{aligned}\mathbb{E}\left[\int_0^\tau \frac{1+\epsilon}{(s+\epsilon)^2(1-s)}\hat{a}_s\, ds\right] &= \mathbb{E}\left[\int_0^\infty \frac{1+\epsilon}{(s+\epsilon)^2(1-s)}\hat{a}_s 1_{\{s\le\tau\}}\, ds\right], \\ &= \int_0^\infty \frac{1+\epsilon}{(s+\epsilon)^2(1-s)}\mathbb{E}\left[\hat{a}_s 1_{\{s\le\tau\}}\right] ds, \\ &= \int_0^\infty \frac{(1+\epsilon)P(\tau\le s)}{(s+\epsilon)^2(1-s)}\mathbb{E}\left[\hat{a}_s|s\le\tau\right] ds.\end{aligned}$$

However, $t \le \tau$ means that $\hat{a}_t \ge g(t)$ for $0 \le s \le t$ and hence $\hat{a}_s|\{t \le \tau\}$ is stochastically larger than $\hat{a}_s$. Thus, $\mathbb{E}\left[\hat{a}_s \mid t \le \tau\right] \ge \mathbb{E}\left[\hat{a}_s\right] = 0$. To complete the proof it is sufficient to use Fatou's lemma:

$$\mathbb{E}\left[\frac{\hat{a}_\tau}{\tau}\right] = \mathbb{E}\left[\liminf_{\epsilon\to 0}\frac{\hat{a}_\tau}{\tau+\epsilon}\right] \le \liminf_{\epsilon\to 0}\mathbb{E}\left[\frac{\hat{a}_\tau}{\tau+\epsilon}\right] \le 0.$$

□

3 FDR Control for Multivariate z-Values

In this section, we extend the step-down BH procedure for multivariate test statistics. The main difficulty in this case arises from the fact that in two, or higher, dimensions we have many choices for enlarging the rejection region. Thus, the original argument based on p-value rankings is not directly applicable. We will prove that the control of FDR remains valid, even when the method for enlarging the rejection region at each level is adapted to z-values rejected prior to that level. This finding can help improve the power of the test—i.e., reduce the FNR—since by observing previously rejected z-values, which likely correspond to non-null hypotheses, we can better estimate the non-null distribution and exploit this information to define the next rejection region more efficiently. An optimal algorithm based on this strategy is presented in Section 4.

3.1 Step-Down BH Procedures in Higher Dimensions

For a given domain $\mathcal{R} \subset \mathbb{R}^d$, we define the following empirical processes

$$\begin{aligned} a(\mathcal{R}) &= \#\{\text{null } z_i : z_i \in \mathcal{R}\}, \\ b(\mathcal{R}) &= \#\{\text{non null } z_i : z_i \in \mathcal{R}\}, \\ r(\mathcal{R}) &= a(\mathcal{R}) + b(\mathcal{R}). \end{aligned}$$

Then,

$$\begin{aligned} \text{FDP}(\mathcal{R}) &= \frac{a(\mathcal{R})}{r(\mathcal{R}) \vee 1}, \\ \text{FDR}(\mathcal{R}) &= \mathbb{E}\left[\frac{a(\mathcal{R})}{r(\mathcal{R}) \vee 1}\right]. \end{aligned}$$

We now present a formal definition of the new step-down procedure.

Definition 1. Let $\mathcal{F}_{\mathcal{R}} = \sigma\left(1_{\{z_i \in B\}} : B \subset \mathcal{R},\ i = 1, \ldots, n\right)$ for all $\mathcal{R} \subset \mathbb{R}^d$. An adaptive step-down BH algorithm is a family of increasing stopping sets $\{\mathcal{R}_t\}_{0 \le t \le 1}$ satisfying the following conditions

1. for $s \le t$ we have $\mathcal{R}_s \subset \mathcal{R}_t$;
2. $F_0(\mathcal{R}_t) = t$;
3. $\forall A \subset \mathbb{R}^d$, $\{\omega : \mathcal{R}_t(\omega) \subset A\} \in \mathcal{F}_A$.

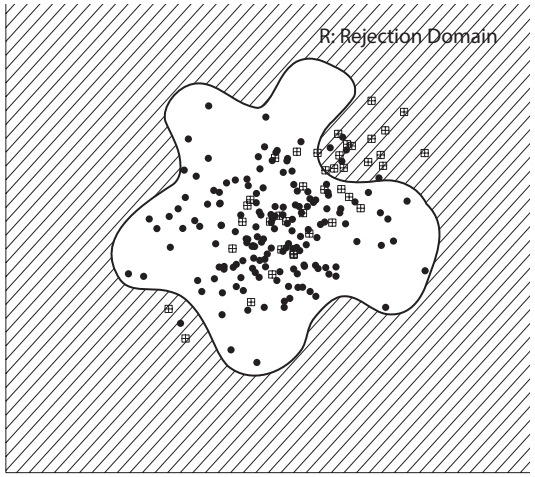

Figure 1: Illustration of adaptive rejection domains; the shaded area $\mathcal{R}$ represents the rejection domain, circles represent null z-values and squares represent alternative hypotheses z-values.

Definition 1 requires that the rejection regions form a family of increasing stopping sets developed based on previously rejected hypotheses. For the family of increasing stopping sets in Definition 1, our proposed generalized step-down BH method reject all z-value in $\mathcal{R}_\tau$, where

$$\tau = \inf \left\{ t : r(\mathcal{R}_t) \leq \frac{nt}{q} \right\}. \tag{8}$$

We next show that, for any error rate q, the proposed generalized BH procedure with the threshold defined in (8) controls the FDR at the level q.

Theorem 3. *If the z-values corresponding to null hypotheses are independent of other z-values, then the generalized step-down BH algorithm controls the false discovery rate at q,*

$$\mathbb{E}\left[\frac{a(\mathcal{R}_\tau)}{r(\mathcal{R}_\tau)}\right] \leq \pi_0 q \leq q. \tag{9}$$

Proof. First note that τ is a stopping time with respect to the filtration $\mathcal{F}_t = \sigma(1_{\{z_i \in \mathcal{R}_s\}}, s \leq t)$. Similar to Section 2, we can rewrite τ as

$$\tau = \inf \Big\{ t : a(\mathcal{R}_t) \leq \underbrace{\frac{nt}{q} - b(\mathcal{R}_t)}_{g(t)} \Big\}.$$

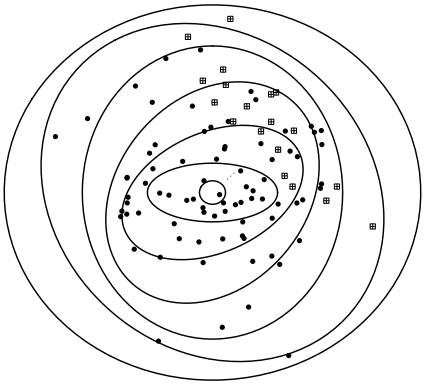

Figure 2: Illustration of the family of increasing nested rejection regions $(\mathcal{R}_t)_{0\le t\le 1}$; the rejection regions are constructed based on previously rejected z-values and are not restricted to have similar shapes.

If we define $\hat{a}(\mathcal{R}_t) = a(\mathcal{R}_t) - n_0 F_0(t)$, then the assertion of Theorem 3 is equivalent to

$$\mathbb{E}\left[\frac{a(\mathcal{R}_\tau)}{r(\tau)}\right] = \mathbb{E}\left[\frac{q}{n}\frac{a(\mathcal{R}_\tau)}{F_0(\tau)}\right] \le \mathbb{E}\left[\frac{q}{n_0}\frac{a(\mathcal{R}_\tau)}{F_0(\tau)}\right] \le q,$$

which is, in turn, equivalent to

$$\mathbb{E}\left[\frac{1}{n_0}\frac{a(\mathcal{R}_\tau)}{F_0(\tau)}\right] \le 1 \Longleftrightarrow \mathbb{E}\left[\frac{\hat{a}(\mathcal{R}_\tau)}{F_0(\tau)}\right] \le 0 \Longleftrightarrow \mathbb{E}\left[\frac{\hat{a}(\mathcal{R}_\tau)}{\tau}\right] \le 0.$$

Now, since the null z-values are independent and $\mathcal{R}_s$ is a stopping set, $a(\mathcal{R}_t) - a(\mathcal{R}_s)$ has a binomial distribution $B\left(n_0 - a(\mathcal{R}_s), \frac{t-s}{1-s}\right)$. Thus,

$$\mathbb{E}\left[a(\mathcal{R}_t)|\mathcal{F}_s\right] = a(\mathcal{R}_s) + (n_0 - a(\mathcal{R}_s))\frac{t-s}{1-s} = \frac{1-t}{1-s}a(\mathcal{R}_s) + \frac{t-s}{1-s}n_0.$$

Given the above equation, and as in the proof of Lemma 2, we conclude that

$$\eta_t = a(\mathcal{R}_t) - \int_0^t \frac{n_0 - a(\mathcal{R}_s)}{1-s}\, ds \tag{10}$$

is an $\mathcal{F}_s$-martingale. The rest of the proof is exactly the same as the proof of Theorem 1. □

4 Oracle Decision Rule for Multivariate Hypothesis Testing

In this section, we construct an oracle decision rule based on *nested rejection regions* for multiple testing in multivariate settings. Inspired by the compound decision rule framework of [21], we first develop an optimal oracle procedure in Theorem 4. Then, in Section 4.1, we propose a nested rejection region algorithm that approximates the oracle rejection rule; we show that our algorithm is asymptotically optimal under weak conditions.

Consider n hypotheses $\mathrm{H}_i, i = 1, \ldots, n$. Denote by θ_i the significance indicator for the i^{th} hypotheses: θ_i is 0 if the i^{th} null hypothesis is true and 1 otherwise. Throughout this section, we assume that θ_1,. . . ,θ_n are independent and identically distributed Bernoulli variables with success probability $\pi_1 = 1 - \pi_0$. The proposed procedure can be generalized to the setting of Markov random fields, as in Sun and Cai [22]. However, such extensions are beyond the scope of the current manuscript and is left to future research. Let $Z = (Z_1, \ldots, Z_n)$ be an independent sequence of test statistics for the n hypotheses H_i with the following density,

$$f(Z_i = z|\theta_i = k) \sim f_k(z) \text{ for } k = 0, 1 \text{ and } i = 1, \ldots, n. \tag{11}$$

Our goal is to construct a multiple testing procedure, based on a realization $z = (z_1, \ldots, z_m)$, which achieves the maximum number of rejected hypothesis, while controlling the FDR at a pre-specified level q. In the next theorem, we assume that the value of π_0 and density function f_1 are known. Under these assumptions, part 1 of Theorem 4 characterizes an optimal set S in terms of Lfdr with optimal power, while controlling the FDR at the level q. Parts 2 and 3 then establish the existence of such a set, whereas part 4 shows that the set S can be obtained by a family of increasing nested stopping sets, defined in Section 3.1. The last part of the theorem provides a specific recipe for constructing optimal rejection region S. Specifically, it shows that the optimal set S can be constructed by applying the BH procedure to the nested regions constructed based on parts 2-4.

Theorem 4. *The following hold for a multiple testing problem with n independent hypotheses represented by d-dimensional test statistics $z_i, i = 1, \ldots, n$ from the mixture distribution* (11).

1. *For a given q if there exists an l such that the set $S = \{x : Lfdr(x) < l\}$ satisfies*

$$\pi_0 \int_S f_0(x)dx = q \int_S f(x)dx,$$

then for any other rejection set T such that $\int_T \pi_0 f_0(x)dx \leq q \int_T f(x)dx$,

$$\int_T f_1(x)dx \leq \int_S f_1(x)dx.$$

2. *If $\inf(Lfdr(x)) > q$, then for any set U with positive Lebesgue measure μ,*

$$\int_U \pi_0 f_0(x)dx > q \int_U f(x)dx.$$

3. *If f is an analytic function and there exists an x such that $Lfdr(x) < q$, then the optimal set S described in 1 exists.*

4. *Under the conditions of part 3, there exists an increasing sequence of nested stopping sets $\{\Omega_t\}_{t=0}^1$ that satisfy the requirements of Definition 1, and for which $\Omega_t = \{x : Lfdr(x) < l_t\}$.*

5. *Let τ_n^{ABH} be the stopping index resulting from applying the adaptive BH procedure of Section 3, i.e.,*

$$\tau_n^{ABH} = \arg\min_t \left\{ q \leq \frac{\pi_0 F_0(\Omega_t)}{\hat{F}(\Omega_t)} \right\}.$$

Further, let τ be the index from part 4 such $S = \Omega_\tau$. Then, as $n \longrightarrow \infty$, $\tau_n^{ABH} \xrightarrow{p} \tau$.

Proof of Theorem 4.

1. The proof of part 1 is similar to proof of Theorem 2.1 in Jin and Zhao [16] with slight modifications.

   First note that since $\text{Lfdr}(x)$ is a monotone decreasing function of the likelihood ratio, $h(x) = \frac{f_1(x)}{f_0(x)}$. Thus, the set S can be obtained by thresholding $h(x)$ at some level l'. Now, for any set T, the definition of S implies that

$$(1_S - 1_T)(f_1(x) - l' f_0(x)) \geq 0.$$

   Let $\alpha(S) = \int_S f_0$ and $\beta(S) = \int_S f_1$ be the type I error and the power of the test based on the rejection region S. Define $\alpha(T)$ and $\beta(T)$ similarly. Then

$$0 \leq \int (1_S - 1_T)(f_1(x) - l' f_0(x))dx = \beta(S) - \beta(T) + l'(\alpha(T) - \alpha(S)). \tag{12}$$

But for any set T satisfying the requirements of the theorem,

$$\pi_0 \int_T f_0(x) - q \int_T f(x)dx \leq 0 = \pi_0 \int_S f_0(x) - q \int_S f(x)dx.$$

Equivalently, $(1-q)\pi_0(\alpha(T) - \alpha(S)) \leq q(1-\pi_0)(\beta(T) - \beta(S))$ or $q'(\alpha(T) - \alpha(S)) \leq \beta(T) - \beta(S)$, where $q' = \frac{(1-q)\pi_0}{q(1-\pi_0)}$. Combining this with Equation (12), we get

$$(l' - q')(\beta(T) - \beta(S)) > 0.$$

Since $\pi_0 \int_S f_0(x)dx = q \int_S f(x)dx$, then $q' \int_S f_0(x) = \int_S f_1(x) < l' \int_S f_0(x)$. Consequently, $l' - q' > 0$, which implies that $\beta(S) \geq \beta(T)$. This completes the proof.

2. If $\inf(\text{Lfdr}(x)) > q$ then for any $x \in U$ we have $\pi_0 f_0(x) - q f(x) > 0$. Now, since U has a positive measure and

$$U = \bigcup_{k \geq 1} U_k, \quad \text{where} \quad U_k = \left\{ x \in U : \pi_0 f_0(x) - q f(x) > \frac{1}{k} \right\},$$

there exists some k for which $\mu(U_k)$ is positive. We then have

$$\int_{U_k} \pi_0 f_0(x) - q f(x)\, dx \geq \int_U \pi_0 f_0(x) - q f(x)\, dx > \int_{U_k} \frac{1}{k}\, d\mu = \frac{\mu(U_k)}{k} > 0,$$

as desired.

3. Let $\Omega_\lambda := \{x : Lfdr(x) < \lambda\}$. We assert that the function $g : \mathbb{R} \longrightarrow \mathbb{R}$, $g(\lambda) := F_0(\Omega_\lambda)$ is continuous. It is sufficient to prove that for any monotone sequence λ_k converging to λ, $g(\lambda_k) \longrightarrow g(\lambda)$. Assume $g(\lambda_k) \nrightarrow g(\lambda)$. We consider two cases. First, assume $\lambda_k \nearrow \lambda$. Then, $\bigcup_{k=1}^{\infty} \Omega_{\lambda_k} \setminus \Omega_\lambda = \emptyset$. Hence $F_0(\Omega_\lambda \setminus \Omega_{\lambda_n}) \longrightarrow 0$, implying that $g(\lambda_n) \longrightarrow g(\lambda)$. Next, suppose $\lambda_k \searrow \lambda$. Then, $\bigcap_{i=1}^{\infty} \Omega_{\lambda_n} \setminus \Omega_\lambda = \Omega^\lambda$, where $\Omega^\lambda := \{x : \text{Lfdr}(x) = \lambda\}$. Therefore, if $g(\lambda_k) \nrightarrow g(\lambda)$ then $F_0(\Omega^\lambda) \neq 0$. But each non-zero empty set has a dense point. Therefore, there exists a sequence $(x_j)_{j=0}^{\infty} \in \Omega^\lambda$, such that $x_j \longrightarrow x$. We next show that all derivatives of Lfdr at point x are zero. First, note that

$$\frac{d}{dx}\text{Lfdr}(x) = \lim_{x_j \to x} \frac{\text{Lfdr}(x_j) - \text{Lfdr}(x)}{x_j - x} = \frac{\lambda - \lambda}{x_j - x} = 0. \qquad (13)$$

Thus, (13) implies that first derivative of Lfdr at x is zero. Higher derivatives can also be computed using the same formula, based only on the values of the function. Thus, all derivatives of Lfdr are zero. However, by the assumption, Lfdr is an analytic function. Now, by the Taylor expansion of Lfdr at point x, we obtain that Lfdr is constant, which is a contradiction. Thus, g is continuous. A similar argument shows that $g(\lambda) := F(\Omega_\lambda)$ is also continuous.

We next characterize the optimal set S. Let $g(\lambda) = \pi_0 \int_{\Omega_\lambda} f_0(x)dx - q \int_{\Omega_\lambda} dF(x)$. Then, $g(\lambda)$ is continuous. Moreover, because $\inf(\text{Lfdr}(x)) > q$, if $\lambda < q$ then $\forall x \in \Omega_\lambda$,

$$\pi_0 f_0(x) < \lambda f(x) \Rightarrow \pi_0 F_0(\Omega_\lambda) < \lambda F(\Omega_\lambda) < qF(\Omega_\lambda) \Rightarrow g(\lambda) < 0.$$

Suppose, without loss of generality, that $\pi_0 > q$. (If $\pi_0 < q$, we can reject all of the hypotheses.) Now, let $\alpha = \sup(\text{Lfdr}(x))$. Then, $g(\alpha) = \pi_0 - q > 0$. Thus, by the mean value theorem, there exists some μ such that $\pi F_0(\Omega_\tau) = qF(\Omega_\tau)$. This completes the proof.

4. Note that $F_0(\Omega_\lambda)$ is a continuous and increasing function with respect to λ. Let $\Omega_0 = \emptyset$. We can then find a function $\alpha : [0, 1] \longrightarrow \mathbb{R}^+$ such that $F_0(\Omega_{\alpha(t)}) = t$. Setting $\mathcal{R}_t = \Omega_{\alpha(t)}$ then gives a family of nested rejection region $\{\mathcal{R}_t\}_{t=0}^1$ satisfying the requirements of Definition 1.

5. Let $\hat{F}_n(\Omega_\lambda) = \frac{1}{n} \sum_{i=1}^{n} I(z_i \in \Omega_\lambda)$ be the empirical distribution function. Then, by Dvoretzky-Kiefer-Wolfowitz inequality (DKW) [11],

$$\mathbb{P}\left(\sup_{t \in [0,1]} |\hat{F}_n(\Omega_\lambda) - F(\Omega_\lambda)| > \epsilon \right) \leq e^{-2n\epsilon^2}.$$

Let $\epsilon_n = n^{-\frac{1}{3}}$. Then, for sufficiently large n, with high probability $F(\Omega_\lambda) - \epsilon_n < \hat{F}_n(\Omega_\lambda) < F(\Omega_\lambda) + \epsilon_n$. Let $\tau_n^{\pm}$ be the solution of equation $\pi_0 F_0(\Omega_t) = qF(\Omega_t) \pm \epsilon_n$. Because the function $g(t) = \dfrac{\pi_0 F_0(\Omega_t)}{qF(\Omega_t)}$ is monotone increasing with respect to t, we have

$$\tau_n^- < \tau_n^{ABH} < \tau_n^+.$$

We assert that τ_n^+ converges to τ. Since $\tau_n^+ \in [0, 1]$ if the sequence diverges, then there exists at least two subsequences $\{\tau_{1,n}^+\}_{n=1}^\infty$ and $\{\tau_{2,n}^+\}_{n=1}^\infty$ converging to, say, τ_1^+ and τ_2^+, respectively. But $g(t)$ is

one to one and $g(\tau_1^+) = g(\tau_2^+) = g(\tau) = q$. Thus, $\tau_1^+ = \tau_2^+ = \tau$, which implies that $\tau^+ \to \tau$. By a similar argument, $\tau^- \to \tau$. Thus, $\tau_n^{ABH} \to \tau$, as desired.

□

Theorem 4 shows that an optimal multiple testing strategy can be obtained using a family of nested rejection regions. Importantly, the theorem establishes that instead of searching over all possible rejection regions in $\mathbb{R}^d$, one only needs to search over the collection of sets Ω_t in order to find optimal rejection set. However, Theorem 4 assumes that f_1 and π_0 are known. We relax this assumption in the next section.

4.1 Approximation of the Oracle Procedure

In Theorem 4, we proved that given the density f_1 of alternative hypotheses and the proportion of null hypotheses π_0, the adaptive BH algorithm of Section 3 converges to the optimal rejection set, as $n \to \infty$. However, in practice, f_1 and π_0 are often unknown. The oracle procedure of Theorem 4 is hence not directly applicable. In Algorithm 1 we propose an approximation to the oracle procedure of Theorem 4 without assuming that f_1 and π_0 are known. For this algorithm, we have:

Theorem 5. *Let $\hat{\Omega}$ be the rejection region obtained from Algorithm 1 for an analytic density function f. Then as $n \to \infty$, $\hat{\Omega}$ converges in probability to some $\Omega_\tau \in \{\Omega_t\}_{t=0}^1$.*

Before presenting the proof of Theorem 5, we comment on two aspects of Algorithm 1, which distinguish it from the oracle procedure of Theorem 4. First, in this algorithm f_1 is assumed to be unknown and is estimated from data. The algorithm builds *nested rejection regions* (NR) by enlarging the rejection region in each step until the stopping criterion is met. Ideally f_1 should be estimated from all data points. In that case, a natural solution would be to estimate f_1, or alternatively f, using a kernel density estimator. However, the construction of adaptive rejection regions in Definition 1 prevents us from using all z-values; instead, we are restricted to using only the rejected z-values at each step of the algorithm. The algorithm thus needs an initial rejection region. To this end, we propose to start by rejecting hypotheses with far enough z-values, e.g. we choose the initial rejection domain to be the complement of a big ball centered at the origin that controls the FDR at a level $q' \ll q$. The control of FDR by such a rejection region

Algorithm 1 The Nested Rejection Region (NR) Algorithm
for Step-Down FDR Control in Multivariate Hypothesis Testing

Initialization:

- Set $\Omega = \Omega_0$ the initial rejection region $\Omega_0 = \{z : \|z\| > \lambda\}$ such that $F_0(\Omega_0) = q' \ll q$.
- Reject z-value in Ω and let $\mathcal{R} = \{z_i : z_i \in \Omega\}$.
- Set FDR $= q'$.

while FDR $\leq q$

1. Define $\hat{f}_\Omega(y) = \dfrac{1}{nh} \sum_{z_i \in \Omega} \omega((y - z_i)/h)$.
2. Reject z^* and add to $\mathcal{R}$ where $z^* = \arg\max_{z_j} \left\{ \frac{\hat{f}_\Omega(z_j)}{f_0(z_i)} : z_j \notin \mathcal{R} \right\}$.
3. Set $\Omega_{\text{search}} = \left\{ z : \frac{\hat{f}_\Omega(z)}{f_0(z)} > \frac{\hat{f}_\Omega(z^*)}{f_0(z^*)} \right\}$.
4. Set $\Omega = \Omega \cup \Omega_{\text{search}}$.
5. Set FDR $= \dfrac{F_0(\Omega)}{n^{-1}\,|\mathcal{R}|}$.

end while

then follows from the validity of the original BH procedure. As the algorithm continues, the rejection region is refined by obtaining a more accurate estimate of f. The main theorem in this section, Theorem 5, shows that if we estimate f by applying a kernel density estimator to z-values in the rejection domain, and use $\hat{f}/f_0$ as the test statistic for extending the rejection domain, then the final rejection region from Algorithm 1 is asymptotically optimal.

The second distinction between Algorithm 1 and the oracle procedure of Theorem 4 concerns the knowledge of the proportion of null hypotheses π_0. In Theorem 4 the optimal rejection region was constructed by assuming that π_0 is known. As a result, the optimal rejection region S in Theorem 4 controls mFDR at the exact level q. However, π_0 is often unavailable. The procedure of Algorithm 1 thus follows the conservative BH approach and (asymptotically) controls mFDR at the level $\pi_0 q$.

Theorem 5 shows that the rejection set obtained from Algorithm 1 approximates the optimal rejection region in Theorem 4, but with mFDR controlled at level $\pi_0 q$. More specifically, Theorem 5 shows that the rejection region from Algorithm 1 converges to $\Omega_\tau \in \{\Omega_\lambda : 0 \leq \lambda \leq 1\}$ introduced in Theorem 4. Each member of this one-parameter family is uniquely characterized by the control level of mFDR. For the rejection region $\hat{\Omega}$ selected by Algorithm 1, we have

$$\lim_{n\to\infty} \frac{F_0(\hat{\Omega})}{n^{-1}\,|\mathcal{R}|} = \pi_0 q.$$

However, Theorem 5 states that, as $n \to \infty$,

$$\frac{F_0(\hat{\Omega})}{n^{-1}\,|\mathcal{R}|} \longrightarrow \frac{F_0(\Omega_\tau)}{F(\Omega_\tau)},$$

for some τ. Thus, Ω_τ satisfies the requirements of Theorem 4, and in particular, satisfies the requirement of part 1 of Theorem 4 at the (conservative) level $\pi_0 q$. This implies that Ω_τ is optimal for FDR control in multivariate hypotheses at the level $\pi_0 q$, which in turn, establishes the (asymptotic) optimality of the rejection region from Algorithm 1.

We next give a proof of Theorem 5.

Proof of Theorem 5. Suppose there exits an initial rejecting set $\mathcal{R}_0$ containing z-values for alternative hypotheses, i.e., $\{z_i : \frac{f(z_i)}{f_0(z_i)} > c_0 \gg 1\} \subset \mathcal{R}_0$. Let $z_{(1)}, z_{(2)}, \ldots, z_{(i),\ldots}$ be sorted z-values according to the step they are rejected in Algorithm 1.

For each i define $\mathcal{R}_i$ to be the rejection region of the i^{th} step containing z-value $z_{(1)}, z_{(2)}, \ldots, z_{(i)}$. For this region, define the estimated density function $\hat{f}_{\mathcal{R}_i}(y) = (ih)^{-1} \sum_{z_k \in \mathcal{R}_i} \omega((y - z_k)/h)$, where ω is an arbitrary kernel. Let $P_i^{(n)} = \min\left\{\frac{f(z_j)}{f_0(z_j)} : z_j \in \mathcal{R}_i\right\}$ and $Q_i^{(n)} = \max\left\{\frac{f(z_j)}{f_0(z_j)} : z_j \notin \mathcal{R}_i\right\}$. If $\mathcal{R}_i$ does not converge to some $\Omega_\lambda = \left\{z : \frac{f(z)}{f_0(z)} > \lambda\right\}$, then for any n there exists an ϵ such that $Q_i^{(n)} - P_i^{(n)} > \epsilon$. We assert that the probability of this event converges to zero for all i, i.e., $\mathbb{P}\left(\max_i\{Q_i^{(n)} - P_i^{(n)}\} > \epsilon\right) \longrightarrow 0 \;\; as\; n \to \infty$.

For each z_0 value, define $\Omega_{\lambda_{z_0}} = \left\{z : \frac{f(z)}{f_0(z)} > \frac{f(z_0)}{f_0(z_0)}\right\}$ and let

$$\tilde{f}_{z_0}(y) = \frac{1}{|\Omega_{\lambda_{z_0}}|h} \sum_{z_k \in \Omega_{\lambda_{z_0}}} \omega((y - z_k)/h).$$

Let $i^* = \text{argmax}_i \left\{ Q_i^{(n)} - P_i^{(n)} \right\}$. Define $P = P_{i^*}^{(n)}$, $Q = Q_{i^*}^{(n)}$ and let z_P and z_Q be the corresponding z-values. For these two z-values we have

$$\frac{\tilde{f}_{z_Q}(z_Q)}{f_0(z_Q)} \geq \frac{\hat{f}_{z_Q}(z_Q)}{f_0(z_Q)} \geq \frac{\hat{f}_{z_P}(z_P)}{f_0(z_P)} \geq \frac{\tilde{f}_{z_P}(z_P)}{f_0(z_P)},$$

where $\tilde{f}_{z_Q}$ is defined similarly as $\tilde{f}_{z_0}$. Therefore,

$$\left\{ \frac{\hat{f}_{z_P}(z_P)}{f_0(z_p)} - \frac{\hat{f}_{z_Q}(z_Q)}{f_0(z_Q)} > \epsilon \right\} \subseteq \left\{ \frac{\tilde{f}_{z_P}(z_P)}{f_0(z_p)} - \frac{\tilde{f}_{z_Q}(z_Q)}{f_0(z_Q)} > \epsilon \right\} \subseteq \left\{ \frac{\tilde{f}_{z_P}(z_P)}{f_0(z_p)} - \frac{\tilde{f}_{z_Q}(z_Q)}{f_0(z_Q)} > 0 \right\}.$$

Let $\Omega_P = \left\{ z : \frac{f(z)}{f_0(z)} > \frac{f(z_P)}{f_0(z_P)} \right\}$ and define Ω_Q similarly. We can write

$$\tilde{f}_{z_P}(x) = \frac{1}{|\Omega_P| h} \sum_{z \in \Omega_P} \omega((x-z)/h) = \frac{1}{|\Omega_P|} \sum_{i=1}^{n} h^{-1} \omega((x-z_i)/h) \mathbb{I}_{\{z_i \in \Omega_P\}},$$

where random variables $h^{-1}\omega((x-z_i)/h)\mathbb{I}_{\{z_i \in \Omega_P\}}$ are i.i.d. To complete the proof, we need the following lemma.

Lemma 6. *The expected value of $\tilde{f}_{z_P}(x)$ at the point z_P converges to $\frac{1}{2} f(z_P)$.*

Proof of Lemma 6. At the point z_P, consider the region $\Omega_P = \left\{ z : \frac{f(z)}{f_0(z)} > \frac{f(z_P)}{f_0(z_P)} \right\}$. Then,

$$\begin{aligned}
\mathbb{E}\left[\tilde{f}_{z_P}(x) \right] &= \mathbb{E}\left[\frac{1}{|\Omega_P|} \sum_{z_i \in \Omega_P} \frac{1}{h} \omega((x-z_i)/h) \mathbb{I}_{\{z_i \in \Omega_P\}} \right] \\
&= \mathbb{E}\left[\frac{1}{h} \omega((x-z_i)/h) \mathbb{I}_{\{z_i \in \Omega_P\}} \right] = \int_{\Omega_P} \omega(u) f(z_P + hu) du.
\end{aligned}$$

For an arbitrary kernel ω, we can choose a $\lambda > 0$ such that $\omega(u) < \frac{\epsilon}{4}$ for $\|u\| > \lambda$. Then, by the Taylor expansion of the analytic density function f,

$$\begin{aligned}
&\int_{\Omega_P} \omega(u) f(z_P + hu) du \\
&\quad = \int_{\Omega_P : \|u\| \leq \lambda} \omega(u) f(z_P + hu) du + \int_{\Omega_P : \|u\| > \lambda} \omega(u) f(z_P + hu) du \\
&\quad < \int_{\Omega_P : \|u\| \leq \lambda} \omega(u) f(z_P + hu) du + \frac{\epsilon}{4}. \qquad (14)
\end{aligned}$$

Let $\bar{\omega}(u) = \omega(u)\mathbb{I}_{\{\|u\|<\lambda\}}$ be the truncated kernel constructed above. By this construction, $\Omega_{P:\|u\|\leq\lambda}$ is within the ball $\mathcal{B}_{h\lambda}(z_P)$ centered at z_P with radius $h\lambda$.

Let $S_P = \left\{z : \frac{f(z)}{f_0(z)} = P\right\}$. Suppose, without loss of generality, that $\nabla\frac{f(z_p)}{f_0(z_p)} \neq 0$ on S_P. Then, S_P has an $(n-1)$-dimensional tangent plane TS_P at z_P, and a one-dimensional orthogonal complement $\langle\nabla\frac{f(z_p)}{f_0(z_p)}\rangle \subset \mathbb{R}^n$. To calculate the integral in (14), we first estimate the region Ω_P. Define the half space $\Omega_T^P = \left\{\mathrm{TS}_P + y\nabla\frac{f(z_p)}{f_0(z_p)} : y \geq 0\right\}$ and let $\Omega_\delta^P = \left\{\mathrm{TS}_P \pm y\nabla\frac{f(z_p)}{f_0(z_p)} : 0 < y < \delta\right\}$. For $\delta = 2h\lambda$, we have $\Omega_T^P \setminus \Omega_\delta^P \subset \Omega_{P:\|u\|\leq\lambda} \subset \Omega_T^P \cup \Omega_\delta^P$. Thus,

$$\left|\int_{\Omega_P:\|u\|\leq\lambda}\omega(u)f(z_P+hu)du - \int_{\Omega_T^P:\|u\|\leq\lambda}\omega(u)f(z_P+hu)du\right| < \int_{\Omega_\delta^P}\omega(u)f(z_P+hu)du.$$

Let M and L be the maximum values of f and ω within the ball $\mathcal{B}_{h\lambda}(z_P)$. Then,

$$\int_{\Omega_\delta^P}\omega(u)f(z_P+hu)du < ML\int_{\mathcal{B}_{h\lambda}(z_P)}du \longrightarrow 0 \text{ as } h \longrightarrow 0.$$

Therefore, we can choose δ_1 such that $\forall h < \delta_1 : \int_{\Omega_\delta^P}\omega(u)f(z_P+hu)du < \frac{\epsilon}{4}$. On the other hand, using a Taylor expansion,

$$\begin{aligned}\int_{\Omega_T^P:\|u\|\leq\lambda}\omega(u)f(z_P+hu)du &= \int_{\Omega_T^P:\|u\|\leq\lambda}\omega(u)[f(z_P)+\mathcal{O}(hu)]du\\ &= f(z_P)\int_{\Omega_T^P}\omega(u)du + \mathcal{O}(h) - \int_{\Omega_T^P:\|u\|>\lambda}\omega(u)f(z_P)du\\ &= \frac{1}{2}f(z_P) - \frac{\epsilon}{4} + \mathcal{O}(h).\end{aligned}$$

In the last equation, we can choose δ_2 such that $\forall h < \delta_2 : \mathcal{O}(h) < \frac{\epsilon}{4}$. Hence for $\forall h < \min\{\delta_1, \delta_2\}$ we have $\left|\int_{\Omega_P}\omega(u)f(z_P+hu)du - \frac{1}{2}f(z_p)\right| < \epsilon$, which completes the proof. □

We now proceed to complete the proof of the theorem. Given that

$\frac{1}{2}\frac{f(z_Q)}{f_0(z_Q)} < \frac{1}{2}\frac{f(z_P)}{f_0(z_p)}$, let $l = \frac{1}{2}\frac{f(z_Q)}{f_0(z_Q)} - \frac{1}{2}\frac{f(z_P)}{f_0(z_p)}$. Then, by Lemma 6,

$$\mathbb{P}\left(\frac{\tilde{f}_{z_Q}(z_Q)}{f_0(z_Q)} - \frac{\tilde{f}_{z_P}(z_P)}{f_0(z_p)} > 0\right) \leq$$
$$\mathbb{P}\left(\left\{\frac{\tilde{f}_{z_Q}(z_Q)}{f_0(z_Q)} > \frac{1}{2}\frac{f_{z_Q}(z_Q)}{f_0(z_Q)} + l/2\right\} \vee \left\{\frac{\tilde{f}_{z_P}(z_P)}{f_0(z_p)} < \frac{1}{2}\frac{f_{z_P}(z_P)}{f_0(z_p)} - l/2\right\}\right) \leq$$
$$\mathbb{P}\left(\frac{\tilde{f}_{z_Q}(z_Q)}{f_0(z_Q)} > \frac{1}{2}\frac{f_{z_Q}(z_Q)}{f_0(z_Q)} + l/2\right) + \mathbb{P}\left(\frac{\tilde{f}_{z_P}(z_P)}{f_0(z_p)} < \frac{1}{2}\frac{f_{z_P}(z_P)}{f_0(z_p)} - l/2\right).$$

We can thus write

$$\mathbb{P}\left(\frac{\tilde{f}_{z_Q}(z_Q)}{f_0(z_Q)} > \frac{1}{2}\frac{f(z_Q)}{f_0(z_Q)} + l/2 \mid Q \text{ is chosen}\right) =$$
$$\mathbb{P}\left(\frac{1}{n}\sum_{i=1}^{n} h^{-1}\omega((x - z_i)/h)\mathbb{I}_{\{z_i \in \Omega_P\}} > \frac{1}{2}f(z_Q) + l_1/2\right), \quad (15)$$

where $l_1 = l f_0(z_Q)$. Now consider a bandwidth h that satisfies $h = \mathcal{O}(1)$; for instance suppose we use the optimal bandwidth $h \sim n^{-\frac{d}{d+4}}$ for d dimensional kernel density estimation. As discussed in the proof of Lemma 6, for every $\epsilon \approx h$ we can find a λ such that for the truncated kernel $\bar{\omega}(u) = \omega(u)\mathbb{I}_{\{\|u\|<\lambda\}}$ in Lemma 6,

$$\mathbb{E}[\bar{\omega}(x - z_i/h)\mathbb{I}_{\{z_i \in \Omega_P\}}] = \frac{1}{2}f(z_Q) + \mathcal{O}(h).$$

Now, let $X_i = \bar{\omega}(x - z_i/h)\mathbb{I}_{\{z_i \in \Omega_P\}}$. Then, by the Hoeffding's inequality [10] we can rewrite Equation (15) as

$$\mathbb{P}\left(\bar{X} - \mathbb{E}[\bar{X}] > l_2/2\right) \leq \exp\left(-\frac{n l_2^2}{2c}\right),$$

where $l_2 = l_1 + \mathcal{O}(h)$. Similarly, we have

$$\mathbb{P}\left(\frac{\tilde{f}_{z_P}(z_P)}{f_0(z_P)} < \frac{1}{2}\frac{f(z_P)}{f_0(z_P)} - l_2/2 \mid P \text{ is chosen}\right) \leq \exp\left(-\frac{n l_2^2}{2c}\right).$$

We thus conclude

$$\mathbb{P}\left(\frac{\tilde{f}_{z_Q}(z_Q)}{f_0(z_Q)} > \frac{\tilde{f}_{z_P}(z_P)}{f_0(z_p)}\right) \leq \frac{1}{2}n(n-1)\exp\left(-\frac{n l_2^2}{2c}\right),$$

where $\frac{1}{2}n(n-1)$ is an upper bound for the number of combinations that P and Q can take. Therefore, as $n \to \infty$, the above probability converges to zero and the proof is complete. □

5 Numerical Experiments

5.1 Simulation Studies

To evaluate the proposed nested region (NR) procedure, we compare its performance with the SC method of Sun and Cai [21], the Fisher's method for combining p-values and an oracle procedure based on known f_1 and π_1. We consider two simulation scenarios:

> *scenario 1*: for each hypothesis, the d-variate test statistics are generated independently;
>
> *scenario 2*: the d-variate test statistics are correlated.

In both settings, the proportion of null hypotheses is set to $\pi_0 = 0.8$ and FDR is controlled at $q = 0.1$ level. For $n \in \{1000, 2000, 5000, 10000\}$ hypotheses, the mean of non-null hypotheses is set to $\mu = 2d^{-1/2}$.

In the first simulation scenario, data corresponding to each hypothesis is generated independently, as $\mathbb{N}(\mu, I_d)$, where $\mu = 0$ for null hypotheses and $\mu = 2d^{-1/2}$ for non-null hypotheses. In other words, in this simulation setting, data corresponding to the d dimensions of each of n hypotheses are independent and the n hypotheses are also independent of each other.

In the second simulation scenario, data for each hypothesis is generated from $\mathbb{N}(\mu, \Sigma_{d\times d})$, where, as before, $\mu = 0$ for null hypotheses and $\mu = 2d^{-1/2}$ for non-null hypotheses. In each simulation instance, the same covariance matrix $\Sigma_{d\times d}$ is used to generate the d-variate test statistics corresponding to the each hypothesis; $\Sigma_{d\times d}$ was generated randomly using the `rcorrmatrix` function in the `clusterGeneration` R-package with parameter `alphad` $= 1$. In other words, in this simulation setting, the multivariate data for each hypothesis are dependent, while the data for two different hypotheses are independent of each other.

To apply NR and SC methods, we use the `pdfCluster` R-package [4] to estimate the density of multivariate test statistics using a kernel of the form

$$\hat{f}(y) = \sum_{i=1}^{n} \left(n \prod_{j=1}^{d} h_j^{opt} \right)^{-1} \prod_{j=1}^{d} K\left(\frac{y_j - x_{i,j}}{h_j^{opt}} \right).$$

In our simulations, the bandwidth for each dimension j is set to the asymptotically optimal bandwidth suggested in [9].

Table 4 and Figure 3 show the estimated FDR and FNR of various methods for simulation scenario 1. The estimated rates are based on averages over over $B = 1000$ simulation replicates. Table 5 and Figure 4 show the same estimates for simulation scenario 2.

In both scenarios, the SC method fails to control the FDR at the desired level, especially as the dimension d increases. In scenario 1, the Fisher's method correctly controls the FDR. This is, of course, expected as the underlying independence assumption of the Fisher's method is satisfied in this case. In simulation scenario 2, where the d-variate test statistics are no longer independent, the Fisher's method fails to control the FDR at the desired level. In fact, the estimated FDR of the Fisher's method is even worst than SC in lower dimensions.

Table 3: The number of rejected hypotheses at various FDR cutoffs using the Fisher's method, the SC method [21] and the proposed NR (nested region) method to the TCGA data. The number of rejected hypotheses using

	FDR Adjustment Methods			BH–mRNA	BH–DNA
q	NR	Fisher	SC	expression	methylation
0.05	52	42	98	16	28
0.1	129	57	167	32	60
0.15	195	86	224	39	101

Unlike the SC and the Fisher's methods, the proposed NR method correctly controls the FDR in both simulation settings and for all values of n and d. The FNR of the NR method is slightly higher than the oracle procedure and the gap increases with d. This is due to the inefficiency of kernel density estimation in higher dimensions. However, the FNR of the NR method improves as the number of hypotheses n increases.

5.2 Application to TCGA Omics Data

To demonstrate the performance of the proposed nested region (NR) approach to detect non-null multivariate hypotheses, we apply it to *omics* data from The Cancer Genome Atlas (TCGA). Specifically, we compare the performance of the NR method with the Fisher's method and the method of Sun and Cai [21] (SC) to control the false discovery rate when testing for "differential activity" of genes with matched DNA methylation and

mRNA expression data from TCGA. To this end, "level 3" expression and methylation data from Colon Adenocarcinoma samples were downloaded from TCGA Data Portal. Matched expression and methylation data for $n = 10676$ genes were retained and separate t-tests were performed based on each omics data type. DNA methylation and mRNA expression capture two different aspect of epigenetic activity of genes in association with cancer. Combining the evidence from these omics data types can help discover novel associations between genes and cancer status from a data integration perspective.

Prior to applying our proposed NR method and the competing approaches, we examined the distribution of the test statistics from expression and methylation data and noticed significant deviations from independence. We thus transformed the test statistics to have identity correlation matrix.

Table 3 shows the number of rejected hypotheses using NR, SC and the Fisher's methods at various FDR cutoffs. As pointed out in Section 5.1, the SC method requires an estimate of the proportion of null hypotheses, π_0. In Section 5.1, we used the true value of π_0 for SC. However, in real data settings, π_0 is unknown and there are currently no methods for estimating π_0 in multivariate settings. We thus estimated π_0 for SC using the smaller of the two estimated values of π_0 for the the methylation and expression data. The estimates of π_0 for each data type were obtained using the lowest slope line approach of Benjamini and Hochberg [6]; this approach resulted in an estimated π_0 of 0.9920 for SC. Similar to the simulation results in Section 5.1, SC rejects the largest number of hypotheses, followed closely by NR and then the Fisher's method. However, as we saw in Section 5.1, there is no guarantee that SC controls the FDR at the desired level. With the transformation applied to the data to obtain an identity correlation matrix, the Fisher's method is expected to control the FDR, as in simulation setting 1 of Section 5.1; however, our simulation results suggest that the Fisher's method may result in high FNR, which may explain the lower number of rejected hypotheses using the Fisher's method.

Table 3 also includes the number of hypotheses rejected if only gene expression or DNA methylation evidence is considered (i.e., results from univariate FDR control). The results of the univariate tests suggest that combining the complementary evidence from the two omics sources results in improved power, highlighting the need for methods to control FDR in multivariate hypotheses. The nested rejection regions from the proposed NR method for this example are shown in Figure 5. Here, the form of the rejection regions are close to circles—which is expected given the transformation to identity correlation matrix—and change slightly as the algorithm

progresses.

6 Discussion

In this paper, a new approach was presented for controlling the false discovery rate (FDR) in multivariate hypothesis testing. In todays increasingly data-driven scientific world, multiple evidences are routinely collected for each hypothesis. Such multivariate hypotheses offer the opportunity for increased statical power and, hence, new scientific discoveries. Unfortunately, existing approaches are designed for univariate hypotheses and are not guaranteed to control the FDR when testing multivariate hypotheses. On the other hand, combining the multivariate evidence into a univariate summary measure (i.e., a single p-value or univariate test statistic) may result in loss of power. The proposed method can thus result in more efficient multivariate hypothesis testing, while controlling the FDR at the desired level.

The approach proposed in this paper is based on a generalization of the proof of the original Benjamini and Hochberg [5] proposal, which allows for more flexible rejection regions defined based on previously rejected hypotheses. The new proof technique is more broadly applicable and can be used to derive more efficient geometric or algorithmic FDR controlling procedures. The nested rejection region (NR) algorithm presented in the paper is an example of such algorithmic approaches, which is shown to be an asymptotically optimal FDR controlling procedure for multivariate hypotheses.

Similar to the original proposal of Benjamini and Hochberg [5], our method assumes that the proportion of null hypotheses π_0 is unknown. It thus controls the FDR conservatively at the level of $\pi_0 q$, instead of the desired FDR level q. The estimation of π_0 for multivariate test statistics is currently an open question and a potentially fruitful research direction. Given such an estimate, more efficient FDR controlling procedures for multivariate hypotheses can be developed. Furthermore, our proposal assumes that both null and non-null hypotheses are independent of each other, while allowing for dependence among d-variate test statistics corresponding to each hypothesis. This assumption is only needed for the development of the optimal algorithm in Section 4.1 and our generalized proof of the BH procedure only requires that the null hypotheses are independent of each other and of non-null hypotheses. Extending the theory in Section 4.1 to allow for dependence among multivariate hypotheses and investigating the effect of dependence among hypotheses can be fruitful areas of future research.

References

[1] ENCODE: ENCyclopedia Of DNA Elements. URL `http://www.genome.gov/ENCODE/`.

[2] The Genotype-Tissue Expression Project (GTEx). URL `http://www.genome.gov/27543767`.

[3] TCGA: The cancer genome atlas: http://cancergenome.nih.gov/. URL `http://cancergenome.nih.gov/`.

[4] A Azzalini, G Menardi, and T Rosolin. R package pdfcluster: Cluster analysis via nonparametric density estimation (version 1.0-0). R-CRAN, 2011.

[5] Y Benjamini and Y Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B. Methodological*, 57(1):289–300, 1995.

[6] Yoav Benjamini and Yosef Hochberg. On the adaptive control of the false discovery rate in multiple testing with independent statistics. *Journal of Educational and Behavioral Statistics*, 25(1):60–83, 2000.

[7] Yoav Benjamini and Daniel Yekutieli. The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics*, 29 (4):1165–1188, 2001.

[8] Yoav Benjamini, Abba M Krieger, and Daniel Yekutieli. Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, 93 (3):491–507, 2006.

[9] Adrian W Bowman and Adelchi Azzalini. *Applied smoothing techniques for data analysis: the kernel approach with S-Plus illustrations: the kernel approach with S-Plus illustrations*. Oxford University Press, 1997.

[10] P. Bühlmann and S. van de Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Series in Statistics. Springer Berlin Heidelberg, 2011. ISBN 9783642201929.

[11] Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics*, pages 642–669, 1956.

[12] Bradley Efron. Size, power and false discovery rates. *The Annals of Statistics*, 35(4):1351–1377, 2007.

[13] Yulia Gavrilov, Yoav Benjamini, and Sanat K Sarkar. An adaptive step-down procedure with proven fdr control under independence. *The Annals of Statistics*, 37(2):619–629, 2009.

[14] Christopher Genovese and Larry Wasserman. A stochastic process approach to false discovery control. *The Annals of Statistics*, 32(3):1035–1061, 2004.

[15] Jiashun Jin and T Tony Cai. Estimating the null and the proportion of nonnull effects in large-scale multiple comparisons. *Journal of the American Statistical Association*, 102(478):495–506, 2007.

[16] Jiashun Jin and Zhigen Zhao. On the generalization of bh procedure. Technical report, 2014.

[17] A. Klenke. *Probability Theory: A Comprehensive Course*. Universitext. Springer London, 2013. ISBN 9781447153610.

[18] Alexander Ploner, Stefano Calza, Arief Gusnanto, and Yudi Pawitan. Multidimensional local false discovery rate for microarray studies. *Bioinformatics*, 22(5):556–565, 2006.

[19] David W Scott. *Multivariate density estimation: theory, practice, and visualization*. John Wiley & Sons, 2015.

[20] John D Storey, Jonathan E Taylor, and David Siegmund. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(1):187–205, 2004.

[21] Wenguang Sun and T Tony Cai. Oracle and adaptive compound decision rules for false discovery rate control. *Journal of the American Statistical Association*, 102(479):901–912, 2007.

[22] Wenguang Sun and T Tony Cai. Large-scale multiple testing under dependence. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(2):393–424, 2009.

[23] Jichun Xie, T Tony Cai, John Maris, and Hongzhe Li. Optimal false discovery rate control for dependent data. *Statistics and its interface*, 4(4):417, 2011.

Table 4: Average FDR and FNR of various multiple comparison adjustment methods in simulation scenario 1; here all $n \in \{1000, 2000, 5000, 10000\}$ (null and non-null) hypotheses are uncorrelated, $q = 0.1$, $\pi_0 = 0.8$, and the mean of non-null hypothesis is set to $2d^{-1/2}$.

		Fisher		SC		NR		Oracle	
n	d	FDR	FNR	FDR	FNR	FDR	FNR	FDR	FNR
1000	2	0.09	0.17	0.14	0.13	0.08	0.15	0.10	0.13
1000	3	0.09	0.18	0.23	0.12	0.09	0.15	0.10	0.13
1000	4	0.10	0.18	0.38	0.11	0.08	0.16	0.10	0.13
1000	5	0.10	0.18	0.53	0.10	0.08	0.16	0.10	0.13
1000	6	0.10	0.18	0.64	0.08	0.08	0.17	0.10	0.12
1000	7	0.09	0.18	0.71	0.07	0.08	0.17	0.10	0.12
1000	8	0.10	0.18	0.76	0.05	0.08	0.18	0.10	0.12
1000	9	0.09	0.19	0.79	0.04	0.09	0.18	0.10	0.12
1000	10	0.10	0.19	0.80	0.02	0.08	0.18	0.10	0.11
2000	2	0.09	0.17	0.12	0.13	0.08	0.14	0.10	0.13
2000	3	0.09	0.18	0.20	0.12	0.08	0.15	0.10	0.13
2000	4	0.09	0.18	0.34	0.11	0.08	0.15	0.10	0.13
2000	5	0.09	0.18	0.49	0.10	0.08	0.15	0.10	0.12
2000	6	0.10	0.18	0.61	0.08	0.08	0.15	0.10	0.12
2000	7	0.10	0.18	0.69	0.07	0.08	0.16	0.10	0.12
2000	8	0.09	0.19	0.75	0.06	0.08	0.16	0.10	0.12
2000	9	0.10	0.19	0.78	0.04	0.08	0.16	0.10	0.12
2000	10	0.09	0.19	0.80	0.03	0.08	0.17	0.10	0.11
5000	2	0.10	0.17	0.11	0.13	0.08	0.14	0.10	0.13
5000	3	0.09	0.18	0.17	0.12	0.08	0.14	0.10	0.13
5000	4	0.09	0.18	0.29	0.12	0.08	0.14	0.10	0.13
5000	5	0.09	0.18	0.44	0.10	0.08	0.14	0.10	0.13
5000	6	0.09	0.18	0.57	0.09	0.08	0.15	0.10	0.12
5000	7	0.09	0.18	0.67	0.07	0.08	0.15	0.10	0.12
5000	8	0.09	0.18	0.73	0.06	0.08	0.15	0.10	0.12
5000	9	0.09	0.19	0.77	0.05	0.08	0.15	0.10	0.12
5000	10	0.09	0.19	0.79	0.03	0.08	0.16	0.10	0.11
10000	2	0.10	0.17	0.11	0.13	0.08	0.14	0.10	0.13
10000	3	0.09	0.18	0.15	0.13	0.09	0.14	0.10	0.13
10000	4	0.09	0.18	0.25	0.12	0.08	0.14	0.10	0.13
10000	5	0.09	0.18	0.40	0.11	0.08	0.14	0.10	0.12
10000	6	0.09	0.18	0.54	0.09	0.08	0.15	0.10	0.12
10000	7	0.10	0.18	0.65	0.08	0.08	0.14	0.10	0.12
10000	8	0.09	0.19	0.71	0.06	0.08	0.15	0.10	0.12
10000	9	0.09	0.19	0.76	0.05	0.08	0.15	0.10	0.12
10000	10	0.10	0.19	0.79	0.04	0.08	0.15	0.10	0.11

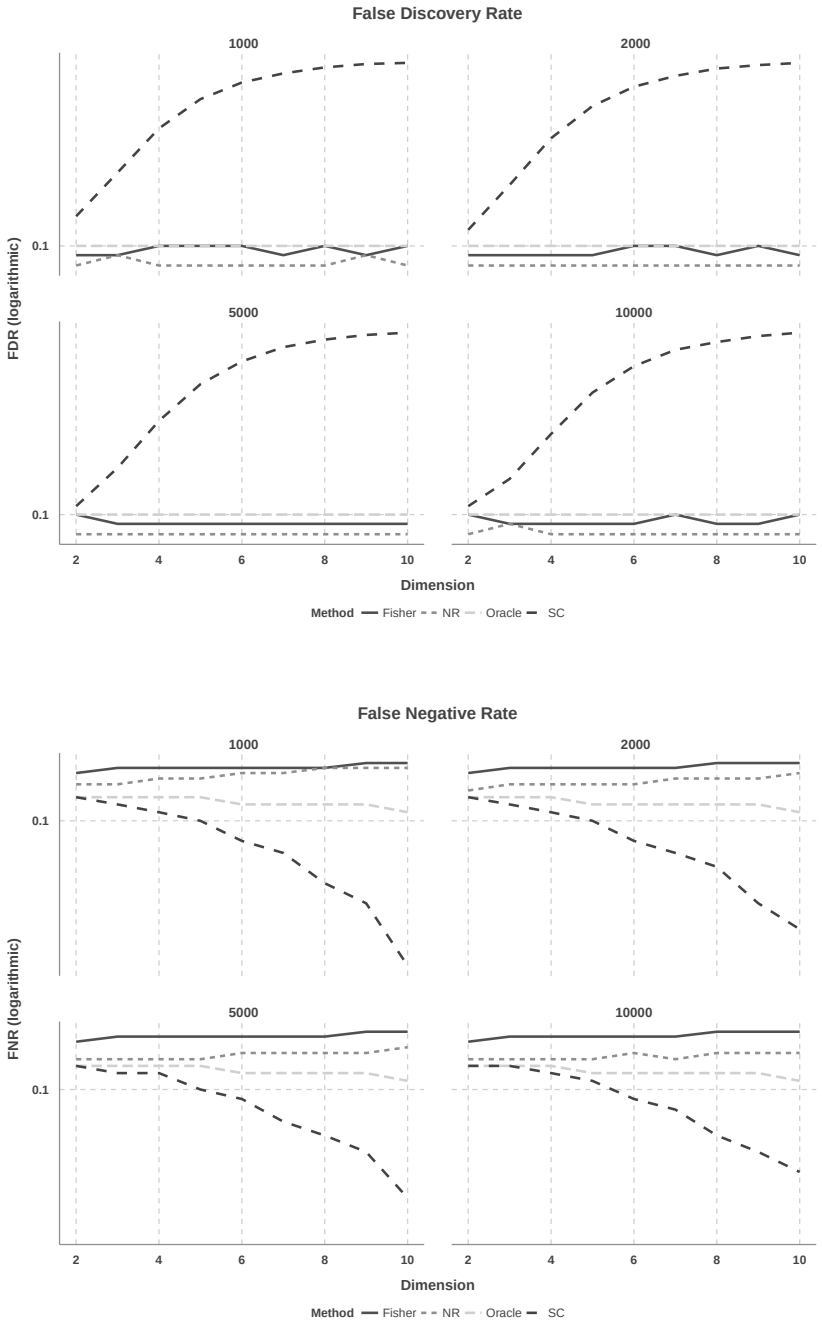

Figure 3: Average FDR and FNR of various multiple comparison adjustment methods for the simulation setting of Table 4.

Table 5: Average FDR and FNR of various multiple comparison adjustment methods in simulation scenario 2; here, $n \in \{1000, 2000, 5000, 10000\}$ d-variate test statistics for each hypothesis are generated from a random correlation matrix, $q = 0.1$, $\pi_0 = 0.8$, and the mean of non-null hypothesis is set to $2d^{-1/2}$.

		Fisher		SC		NR		Oracle	
n	d	FDR	FNR	FDR	FNR	FDR	FNR	FDR	FNR
1000	2	0.33	0.13	0.13	0.13	0.08	0.16	0.10	0.13
1000	3	0.37	0.15	0.22	0.12	0.08	0.15	0.10	0.11
1000	4	0.40	0.15	0.38	0.10	0.08	0.15	0.10	0.09
1000	5	0.42	0.16	0.53	0.09	0.08	0.15	0.10	0.07
1000	6	0.43	0.16	0.65	0.08	0.09	0.15	0.10	0.05
1000	7	0.44	0.16	0.72	0.07	0.09	0.15	0.10	0.03
1000	8	0.45	0.17	0.77	0.07	0.08	0.15	0.10	0.02
1000	9	0.46	0.17	0.79	0.06	0.09	0.15	0.10	0.02
1000	10	0.46	0.17	0.80	0.04	0.08	0.16	0.11	0.01
2000	2	0.33	0.13	0.12	0.13	0.08	0.16	0.10	0.13
2000	3	0.37	0.15	0.19	0.12	0.08	0.15	0.10	0.11
2000	4	0.39	0.15	0.33	0.10	0.08	0.14	0.11	0.09
2000	5	0.41	0.16	0.49	0.09	0.08	0.14	0.10	0.07
2000	6	0.43	0.16	0.62	0.08	0.08	0.13	0.10	0.05
2000	7	0.44	0.16	0.70	0.07	0.08	0.13	0.10	0.04
2000	8	0.45	0.17	0.76	0.07	0.08	0.13	0.10	0.02
2000	9	0.46	0.17	0.79	0.06	0.08	0.13	0.10	0.02
2000	10	0.46	0.17	0.80	0.03	0.08	0.13	0.11	0.01
5000	2	0.33	0.13	0.11	0.13	0.07	0.16	0.10	0.13
5000	3	0.37	0.15	0.16	0.12	0.07	0.15	0.10	0.11
5000	4	0.39	0.15	0.27	0.10	0.07	0.14	0.10	0.09
5000	5	0.41	0.16	0.43	0.09	0.08	0.13	0.10	0.07
5000	6	0.43	0.16	0.57	0.08	0.07	0.13	0.10	0.05
5000	7	0.44	0.16	0.68	0.07	0.07	0.12	0.11	0.04
5000	8	0.45	0.17	0.74	0.07	0.07	0.12	0.10	0.02
5000	9	0.46	0.17	0.78	0.06	0.07	0.13	0.10	0.01
5000	10	0.46	0.17	0.80	0.05	0.07	0.12	0.11	0.01
10000	2	0.33	0.13	0.11	0.13	0.08	0.15	0.10	0.13
10000	3	0.37	0.15	0.14	0.12	0.07	0.14	0.10	0.11
10000	4	0.40	0.15	0.24	0.10	0.07	0.14	0.10	0.09
10000	5	0.41	0.16	0.39	0.09	0.07	0.13	0.10	0.07
10000	6	0.43	0.16	0.54	0.08	0.07	0.12	0.10	0.05
10000	7	0.44	0.16	0.65	0.07	0.07	0.12	0.10	0.04
10000	8	0.45	0.17	0.73	0.06	0.07	0.11	0.11	0.02
10000	9	0.46	0.17	0.77	0.06	0.07	0.11	0.11	0.02
10000	10	0.46	0.17	0.79	0.05	0.06	0.11	0.11	0.01

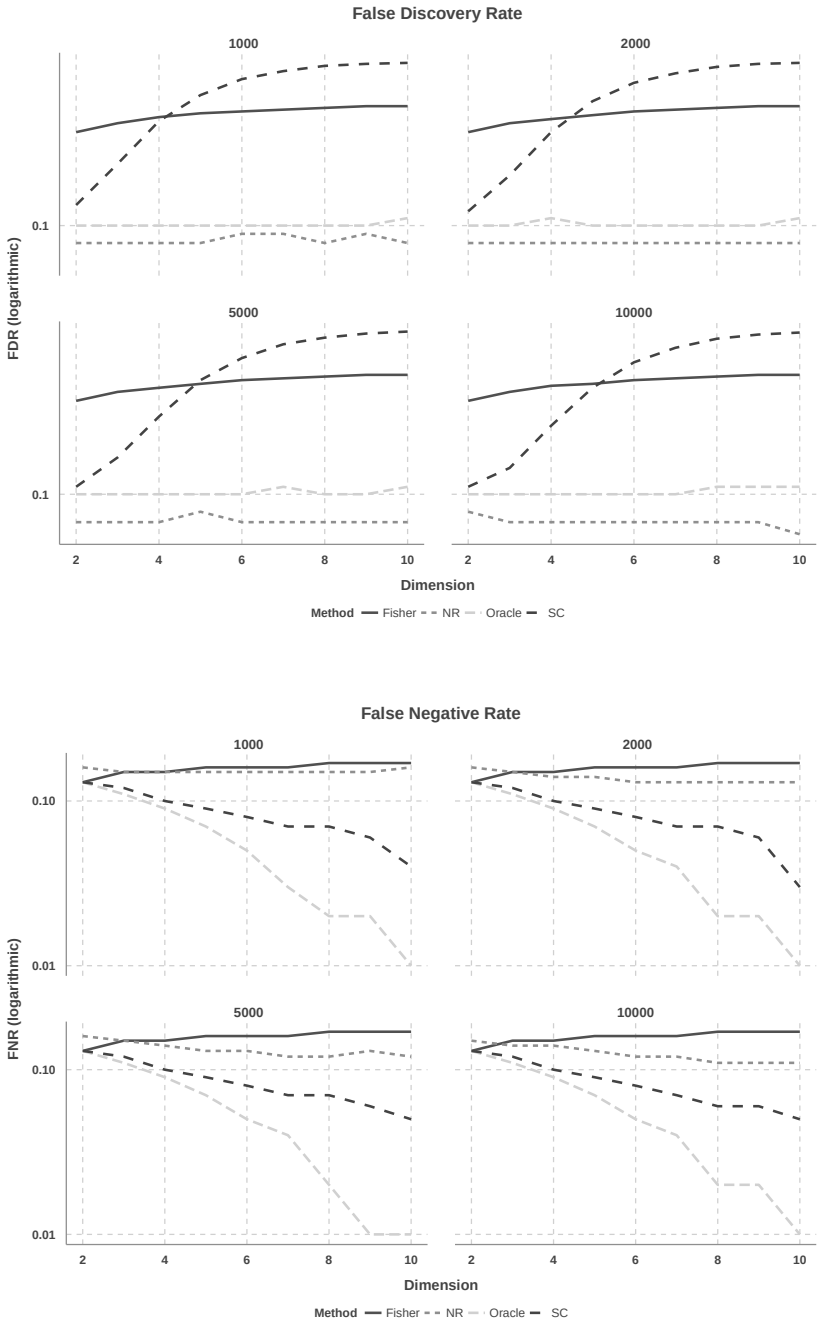

Figure 4: Average FDR and FNR of various multiple comparison adjustment methods for the simulation setting of Table 5.

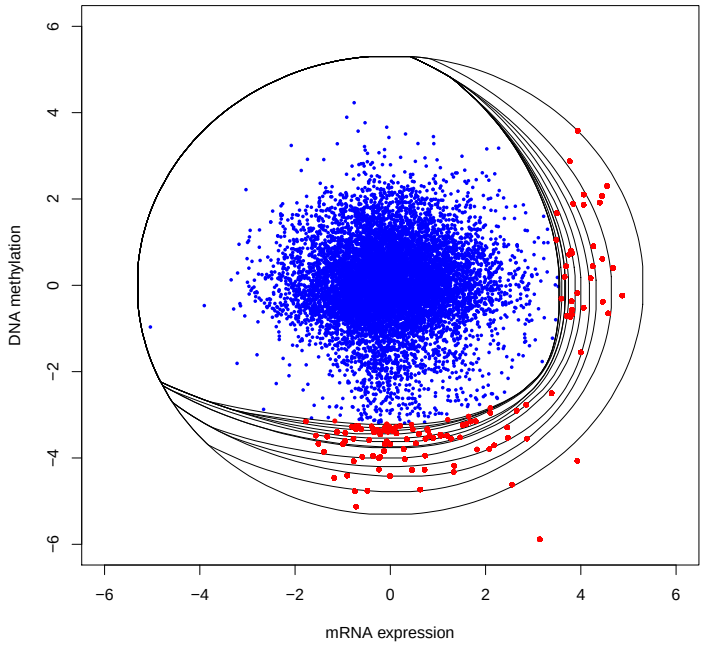

Figure 5: Nested rejection regions for test statistics from the TCGA omics data; small blue circles represent the final null hypotheses, while bigger red points represent non-null hypotheses.